

Specular Highlights Detection and Reduction with Multi-Flash Photography

Rogério Feris¹ Ramesh Raskar² Kar-Han Tan³ Matthew Turk¹

¹University of California, Santa Barbara, USA
{rferis,mturk}@cs.ucsb.edu

²Mitsubishi Electric Research Labs, USA
raskar@merl.com

³Epson Palo Alto Research Labs, USA
tankh@vision.ai.uiuc.edu

Abstract

We present a novel method to reduce the effect of specularities in digital images. Our approach relies on a simple modification of the capture setup: a multi-flash camera is used to take multiple pictures of the scene, each one with a differently positioned light source. We then formulate the problem of specular highlights reduction as solving a Poisson equation on a gradient field obtained from the input images. The obtained specular reduced image is further refined in a matting process with the maximum composite of the input images. Experimental results are demonstrated on real and synthetic images. The entire setup can be conceivably packaged into a self-contained device, no larger than existing digital cameras.

Keywords: Specularities, Multi-Flash Imaging, Poisson Equation

Correspondence Author:

Rogério Feris
6661 Berkshire Terrace apt. 18
Goleta, CA 93117
USA
(805)284-5392
rferis@cs.ucsb.edu

1 Introduction

The reflection of light from surfaces in real scenes is generally classified into two main categories: diffuse and specular. The diffuse component results from light rays penetrating the surface, undergoing multiple reflections, and re-emerging [9]. In contrast, the specular component is a surface phenomenon - light rays incident on the surface are reflected such that the angle of reflection equals the angle of incidence. Light energy due to specular reflections is often concentrated in a compact lobe, causing strong highlights (bright regions) to appear in the image.

These bright spots, also known as specularities, play a major role in many computer graphics and vision problems. They provide a true sense of realism in the environment, reveal important local curvature information [10] and may even provide additional cues for object recognition [11]. However, in most cases, specular highlights are undesirable in images. They are often considered as annoyance in traditional photography and cause vision algorithms for segmentation and shading analysis to produce erroneous results. If the sensor direction is varied, highlights shift, diminish rapidly, or suddenly appear in other parts of the scene. This poses a serious problem for vision methods that rely on image correspondences, such as stereo or motion algorithms.

A variety of photometric techniques have been proposed for detecting and removing specularities using color [7], polarization [17], multiple views [8] and hybrid methods [9]. However, most of these techniques assume that highlights appear in regions with no variation of material type or surface normal. In fact, reliably removing specularities in textured regions remains a challenging problem.

In this paper, we address the problem of reducing the effect of specularities in digital images by taking a different approach¹. Rather than relying in one single image as most previous methods, we capture a set of images of the scene with different lighting conditions. Our work belongs to an emerging class of computer graphics techniques that process multiple images acquired with the same viewpoint but under different conditions, such as different illumination, focus or exposure [3, 1, 13].

¹Based on “Specular Reflection Reduction with Multi-Flash Imaging”, by Feris, Raskar, Tan, and Turk, which appeared in Proceedings of SIBGRAPI/SIACG 2004 (ISBN 0-77-695-2227-0). ©2004 IEEE.




Figure 1: *Multi-flash camera setups used in [13] for depth edge detection. From left to right: 4-flash setup, 8-flash setup and dynamic scenes setup. Our goal in this paper is to exploit multi-flash imaging for specular reflection reduction.*

Our capture setup is a multi-flash camera, a self-contained device that we have recently proposed for depth edge extraction and non-photorealistic rendering [13]. By observing that specularities shift according to the different positions of the flashes, we are able to significantly reduce or completely remove specular highlights in the scene. More specifically, we formulate this problem as solving a Poisson equation on a gradient field obtained from the input images. This allows us to remove specularities in objects with different curvatures and textured regions.

We will detail our approach in Section 2 and present our experimental results in Section 3. In Section 4, we discuss advantages and limitations of our method. In Section 5, we describe implementation details and conclude the paper in Section 6, with final remarks and future work.

2 Specular Highlights Reduction

In this section we introduce our proposed method for specular highlights reduction. We start describing our capture setup and then present our gradient domain algorithm.

2.1 Multi-Flash Imaging

Instead of taking one single picture of the scene, we use a multi-flash camera with n flashes to acquire n images from the same viewpoint, each one with a differently positioned flash. We have already successfully used this setup for detection of depth

edges, with applications in stylized rendering [13], fingerspelling recognition [5] and medical imaging [14]. Figure 1 shows our different multi-flash setups for static and dynamic scenes. Note that the position of each flash in these prototypes was chosen to better detect depth edges. We refer to [13] for details about our depth edge detection method based on shadows. Now we will show how multi-flash cameras can be used for reliable removal and reduction of specularities.

2.2 Approach

Our method is based on the observation that specular spots shift according to the shifting of light sources that created them. We need to consider **three cases** of how specular spots in different light positions appear in each image:

- (i) shiny spots remain distinct on a specular surface.
- (ii) some spots overlap.
- (iii) spots overlap completely (no shift).

We show that for cases (i) and (ii), which often occur in practice, our method successfully removes specular highlights.

We note that although specularities overlap in the input images, the boundaries (intensity edges) around specularities in general do not overlap. The main idea is to exploit the gradient variation in the n images, taken under the n different lighting conditions, at a given pixel location (x,y) . If (x,y) is in a specular region, in cases (i) and (ii), the gradient due to the specular boundary will be high in only one or a minority of the n images. Taking the **median of the n gradients** at that pixel will remove this outlier(s). Our method is motivated by the intrinsic image approach [16], where the author removes shadows in outdoor scenes by noting that shadow boundaries are not static.

Let I_k , $1 \leq k \leq n$ be an input image taken with light source k and $Imax(x) = max_k(I_k(x))$ be the maximum composite of the input images. We compute the specular-reduced image through the following steps:

1. Compute intensity gradient, $G_k(x, y) = \nabla I_k(x, y)$
2. Find median of gradients, $G(x, y) = median_k(G_k(x, y))$




Figure 2: The basic idea of our algorithm for specular reflection reduction.

3. Reconstruct image \hat{I} which minimizes $|\nabla \hat{I} - G|$
4. Use $Imax$ to improve result in a matting process.

This algorithm is illustrated in Figure 2, considering a camera with four flashes. We will detail the method used for reconstructing image \hat{I} (step 3) in Section 2.3. The alpha matting process (step 4) will be described in Section 2.4.

Figure 3 shows a simple example to illustrate our method in all three cases mentioned above. For each case, we created four images with manually drawn specularities. The first column in the figure corresponds to the max composite of the four images (I_{max}), the second corresponds to the median of intensities (I_{median}) and the third column is the output of our method - the reconstruction from the median of gradients ($I_{intrinsic}$).

Note that if we consider I_{median} , specularities are not eliminated in case (ii), where spots overlap. On the other hand, our method is able to handle cases (i) and (ii), which often occur in practice. If specularities do not move among images, our method fails to remove them.

It is worth mentioning that specularities could also be removed by just taking the min composite of the images, but in this case we would have the presence of shadows, which are not desirable. Since the boundaries of shadows rarely overlap, they are also eliminated in our method.

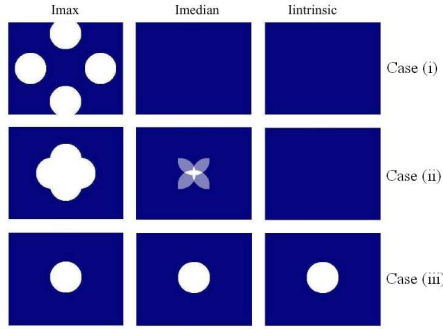


Figure 3: *Illustration of the three cases. Note that if we consider only the median of image intensities (instead of median of gradients), we have problems in case (ii). Our method based on the intrinsic image handles cases (i) and (ii) which often occur in practice. If specularities do not move among images our method fails to remove them.*

2.3 Image Reconstruction from Gradient Fields

Image reconstruction from gradients fields, an approximate invertibility problem, is still a very active research area. In R^2 , a modified gradient vector field G may not be integrable. In other words, there might not exist an image \hat{I} such that $G = \nabla \hat{I}$. In fact, the gradient of a potential function must be a conservative field i.e., the gradient $\nabla \hat{I} = (\frac{\partial \hat{I}}{\partial x}, \frac{\partial \hat{I}}{\partial y})$ must satisfy:

$$\frac{\partial^2 \hat{I}}{\partial x \partial y} = \frac{\partial^2 \hat{I}}{\partial y \partial x} \quad (1)$$

This condition is rarely verified for the gradient field G .

As noted by Frankot and Chellappa [6], one possible solution to this problem is to orthogonally project G onto a finite set of Fourier basis functions spanning the set of integrable vector fields. In our method, we employ a more direct and efficient approach, in a similar way to the work of Fattal et al. [4]. We search the space of all 2D potential functions for a function \hat{I} whose gradient is the closest to G in the least-squares sense. In other words, \hat{I} should minimize the integral:

$$\int \int F(\nabla \hat{I}, G) dx dy, \quad (2)$$

where $F(\nabla \hat{I}, G) = \|\nabla \hat{I} - G\|^2 = (\frac{\partial \hat{I}}{\partial x} - G_x)^2 + (\frac{\partial \hat{I}}{\partial y} - G_y)^2$.

According to the Variational Principle, a function \hat{I} that minimizes the integral in

(2) must satisfy the Euler-Lagrange equation:

$$\frac{\partial F}{\partial \hat{I}} - \frac{d}{dx} \frac{\partial F}{\partial \hat{I}_x} - \frac{d}{dy} \frac{\partial F}{\partial \hat{I}_y} = 0, \quad (3)$$

which is a partial differential equation in \hat{I} . Substituting F we obtain the following equation:

$$2\left(\frac{\partial^2 \hat{I}}{\partial x^2} - \frac{\partial G_x}{\partial x}\right) + 2\left(\frac{\partial^2 \hat{I}}{\partial y^2} - \frac{\partial G_y}{\partial y}\right) = 0 \quad (4)$$

Dividing by 2 and rearranging terms, we obtain the well-known Poisson Equation:

$$\nabla^2 \hat{I} = \text{div}G \quad (5)$$

where ∇^2 is the Laplacian operator defined as $\nabla^2 \hat{I} = \frac{\partial^2 \hat{I}}{\partial x^2} + \frac{\partial^2 \hat{I}}{\partial y^2}$, and $\text{div}G$ is the divergence of the vector field G , defined as $\text{div}G = \frac{\partial G_x}{\partial x} + \frac{\partial G_y}{\partial y}$.

We have used the standard full multigrid method [12] to solve the Poisson equation. We pad the images to square images of size the nearest power of two, before applying the integration, and then crop back the result to the original size.

2.4 Alpha Matting

As we will show in our experiments, our technique is effective to reduce the effect of specularities in digital images, but may also introduce some low-contrast effects, which might be undesirable for some applications (e.g., photography).

We improve the quality of our result in a matting process with the maximum composite image I_{max} . We place the flashes strategically so that every point in the scene that is shadowed in some image is also captured without being shadowed in at least one other image. Thus, I_{max} is a shadow-free image, with the problem of being affected by specularities.

Our observation is that we can *detect* specularities in I_{max} by taking the difference between I_{max} and the reconstructed specular-reduced image \hat{I} . This difference is then used as an alpha channel for matting, so that we replace the specular regions in I_{max} by the correspondent specular-reduced regions in \hat{I} . More specifically, we computed

our final result R as follows:

$$\begin{aligned}\alpha &= \frac{Imax - \hat{I}}{N} \\ R &= \alpha \hat{I} + (1 - \alpha)Imax\end{aligned}\tag{6}$$

where N is a normalizing constant factor, calculated so that values of α range from zero to one .

An alternative to the matting process would be to threshold α and use image inpainting techniques [2, 15] to fill in the specular regions in $Imax$ (determined by α).

3 Results

This section presents our experimental results, which were carried out with synthetic and real images, considering objects with different local curvature and textured regions.

3.1 Synthetic Images

We first demonstrate our method in synthetic images, showing that it can reliably remove specularities in textured regions, as long as spots shift among images.

Figure 4a-d shows four images of a textured region, with manually drawn specularities (white circles). Figure 4e shows the max composite image (case (ii), where spots overlap) and our result is shown in Figure 4f. Note that we are able to eliminate the specularities, while preserving the texture.

3.2 Real Images

We now show the usefulness of our method in real scenes. Figure 5a-d shows four images taken with our multi-flash setup depicted in Figure 1a. Note that the images contain shadows and specularities. The max composite image, shown in Figure 5e, has no shadows, but contains all specular highlights. Our specular-reduced gradient reconstruction is shown in Figure 5f. We have chosen objects in the scene with different




Figure 4: (a-d) Four images with manually drawn specularities along a textured region. (e) Max composite image. (f) Result of our method.

local curvature and included also a textured box. We are able to significantly remove specular highlights despite these conditions.

Figure 5g shows the magnitudes of gradients along the scanline showed in the max composite image. Dashed lines correspond to gradients of different images and the solid line is the median of gradient magnitudes. Figure 5h shows the intensity of the max composite image along the scanline (dashed line) and the intensity of the reconstructed specular-reduced image (solid line).

The quality of our result is improved through the alpha matting step. Figure 5i shows the computed alpha channel, which allows us to detect specularities. Our final result after matting is shown in Figure 5j.

Figure 6a shows a challenging scene, involving a transparent object and a high albedo background. Note that we are able to detect most specularities (Figure 6c) without false positives. As we can see in our final result for this example (Figure 6d), shadows are eliminated and specularities are significantly reduced.

Another demonstration of our technique is given in Figure 7, where we show an input image taken with one of the flashes (Figure 7a), the gradient reconstructed image (Figure 7b), and our final result after matting (Figure 7c). We note in this case that the gradient reconstructed image has some blur and shadow effects (due to overlapped shadow boundaries). The quality of our result is improved after alpha matting.

The computational time required to obtain the specular-free images is about one second on a 3GHz Pentium 4, considering images with resolution 640x480.




Figure 5: (a-d) Four images taken with our multi-flash camera. (e) Max composite image. (f) Specular-reduced image using our method. (g) Magnitudes of gradients along the scanline showed in (e). Dashed lines correspond to gradients of different images and the solid line is the median of gradient magnitudes. (h) Intensity of max composite along the scanline (dashed line) and intensity of reconstructed specular-reduced image (solid line). (i) Alpha channel used for matting. (j) Our final result.

4 Discussion

As we already mentioned, our method fails when specular highlights do not shift among images captured with different flashes. Fortunately, this case is not problematic for many computer vision methods that rely on image correspondences. Other sources of failure include scenes with e.g., rough surfaces or silk cloths.

In our method, regions covered by specular highlights in one image may be specular-free in others, since bright spots often shift due to our multi-flash scheme. This allows us to “uncover” these regions, posing an advantage over previous approaches that rely on a single image.

By solving the Poisson partial differential equation, we eliminate most highlights




Figure 6: (a). Image taken with one of the flashes. (b). Maximum composite image. (c). Detection of specularities (alpha channel). (d) Our final result after matting.




Figure 7: (a) Image taken with one of the flashes. (b) Gradient reconstructed image. (c) Result after matting.

and also shadows in the image. This compares favorably with the max composite (which has no shadows but is full of specularities), the min composite (which is specular-reduced but also full of shadows) and the median of intensities (which has problems when spots overlap, as shown in Figure 3). In addition, our method is not affected by objects as bright as specularities in the scene (see Figure 6, which is clearly a problem for thresholding-based techniques).

We have tested our approach using the same setup for depth edge detection [13] with four flashes. However, even more accurate results could be obtained by using more light sources with strategic positions. Using only two flashes would not be a good choice, due to the fact that the gradient of the boundary of a specularity in one image would be only attenuated (not removed) when taking the mean with the other image. Choosing the minimum gradient could help in this case, but it would create a problem for textured regions.

In addition to reduce the effect of specular highlights in digital images, our method




Figure 8: (a) Max composite image and (b) correspondent depth edges. (c) Intensity plots along the red scanline of the region highlighted in leftmost figure. (d) Specular mask computation (e) Our final result

also detects specularities, which may be useful in other applications, such as object shape information acquisition [9] and interactive creation of technical illustrations, where the user may decide whether to keep or remove specularities. Our detected specular regions could also be used as input for image inpainting methods, which often require manual labeling of such regions [15].

Our technique could be applied for specular reflection reduction in traditional photography and also for improving computer vision methods. We show in Figure 8 that using our specular-reduced image significantly improves depth edge detection [13]. Figure 9 shows more results in challenging specular scenes.

Finally, we need to mention that we are processing “active specularities”, created by our light sources, rather than highlights due to ambient light. The use of flashes might be useful in many situations, such as in dark scenes, or even for vision methods based on active lighting.

5 Implementation

Our basic prototype for static scenes makes use of a four MegaPixel Canon Powershot G3 digital camera. The four booster (slaved Quantarray MS-1) 4ms duration flashes are triggered by optically coupled LEDs turned on sequentially by a PIC microcontroller, which in turn is interrupted by the hot-shoe of the camera.




Figure 9: *Depth edge detection results in challenging specular scenes.*

6 Conclusions

We have presented a new method for specular reflection reduction in digital images. Rather than relying on a single image, we capture multiple images of the scene under different lighting conditions. Based on the information that highlights often shift according to the position of the light sources, we formulate the problem of specular-ity removal as solving a Poisson equation on a gradient field obtained from the input images.

We basically improved and presented another functionality for the multi-flash camera, recently proposed for depth edge detection and stylized rendering [13]. As future work, we plan to exploit other imaging parameters, such as variable wavelength (coloured flashes) to improve our techniques.

References

- [1] D. Akers, F. Losasso, J. Klingner, M. Agrawala, J. Rick, and P. Hanrahan. Conveying Shape and Features with Image-Based Relighting. In *IEEE Visualization*, 2003.
- [2] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester. Image inpainting. In *Siggraph 2000, Computer Graphics Proceedings*, pages 417–424. ACM Press / ACM SIGGRAPH / Addison Wesley Longman, 2000.
- [3] M. Cohen, A. Colburn, and S. Drucker. Image stacks. Technical Report MSR-TR-2003-40, Microsoft Research, 2003.

- [4] R. Fattal, D. Lischinski, and M. Werman. Gradient Domain High Dynamic Range Compression. In *Proceedings of SIGGRAPH 2002*, pages 249–256. ACM SIGGRAPH, 2002.
- [5] R. Feris, M. Turk, R. Raskar, K. Tan, and G. Ohashi. Exploiting Depth Discontinuities for Vision-based Fingerspelling Recognition. In *IEEE Workshop on Real-time Vision for Human-Computer Interaction (in conjunction with CVPR'04)*, Washington DC, USA, 2004.
- [6] R. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451, 1988.
- [7] G. Klinker, S. Shafer, and T. Kanade. The measurement of highlights in color images. *International Journal of Computer Vision*, 2:7–32, 1988.
- [8] S. Lee. *Understanding of Surface Reflections in Computer Vision by Color and Multiple Views*. PhD thesis, University of Pennsylvania, 1991.
- [9] S. Nayar, X. Fang, and T. Boult. Removal of specularities using color and polarization. In *International Conference on Computer Vision and Pattern Recognition*, pages 583–590, New York City, USA, 1993.
- [10] M. Oren and S. Nayar. A theory of specular surface geometry. *International Journal of Computer Vision*, 2:105–124, 1997.
- [11] M. Osadchy, D. Jacobs, and R. Ramamoorthi. Using specularities for recognition. In *International Conference on Computer Vision*, Nice, France, 2003.
- [12] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. Pearson Education, 1992.
- [13] R. Raskar, K. Tan, R. Feris, J. Yu, and M. Turk. A non-photorealistic camera: depth edge detection and stylized rendering using multi-flash imaging. *SIGGRAPH'04 / ACM Transactions on Graphics (to appear)*, 2004.

- [14] K. Tan, J. Kobler, P. Dietz, R. Feris, and R. Raskar. Shape-Enhanced Surgical Visualizations and Medical Illustrations with Multi-Flash Imaging. In *Conference on Medical Image Computing and Computer Assisted Intervention*, France, 2004.
- [15] P. Tan, S. Lin, L. Quan, and H. Shum. Highlight removal by illumination-constrained inpainting. In *International Conference on Computer Vision (ICCV'03)*, pages 164–169, 2003.
- [16] Y. Weiss. Deriving intrinsic images from image sequences. In *Proceedings of International Conference on Computer Vision*, pages 68–75, 2001.
- [17] L. Wolff and T. Boult. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(6):635–657, 1991.